REAL-TIME ENSEMBLE CONTROL WITH REDUCED-ORDER MODELING

BINGHUAI LIN* AND DENNIS MCLAUGHLIN†

Abstract. The control of spatially distributed systems is often complicated by significant uncertainty about system inputs, both time-varying exogenous inputs and time-invariant parameters. Spatial variations of uncertain parameters can be particularly problematic in geoscience applications, making it difficult to forecast the impact of proposed controls. One of the most effective ways to deal with uncertainties in control problems is to incorporate periodic measurements of the system's states into the control process. Stochastic control provides a convenient way to do this, by integrating uncertainty, monitoring, forecasting, and control in a consistent analytical framework. This paper describes an ensemble-based approach to closed-loop stochastic control that relies on a computationally efficient reduced-order model. The use of ensembles of uncertain parameters and states makes it possible to consider a range of probabilistic performance objectives and to derive real-time controls that explicitly account for uncertainty. The process divides naturally into measurement updating, control, and forecasting steps carried out recursively and initialized with a prior ensemble that describes parameter uncertainty. The performance of the ensemble controller is investigated here with a numerical experiment based on a solute transport control problem. This experiment evaluates the performance of open and closed-loop controllers with full and reduced-order models as well as the performance obtained with a controller based on perfect knowledge of the system and the nominal performance obtained with no control. The experimental results show that a closed-loop controller that relies on measurements consistently performs better than an open loop controller that does not. They also show that a reduced-order forecasting model based on off-line simulations gives nearly the same performance as a significantly more computationally demanding full order model. Finally, the experiment indicates that a moderate penalty on the variance of control cost yields a robust control strategy that reduces uncertainty about system performance with little or no increase in average cost. Taken together, these results confirm that reduced-order ensemble closed-loop control is a flexible and efficient control option for uncertain spatially distributed systems.

Key words. model order reduction, ensemble Kalman filter, real-time control, model predictive control

AMS subject classifications. 93E11, 93E20, 93B11

1. Introduction.

Efforts to control dynamic systems can often benefit from frequent measurements of system behavior. Such measurements can help compensate for inaccurate predictions or other errors by revealing the actual systems response to control actions. If we have a perfect model of the system's response there is no need to rely on measurements since they do not add any new information. But in realistic situations, where models are imperfect and system response is uncertain, proper use of measurement information can greatly improve system performance. The roles of uncertainty and measurement feedback are closely related and deserve to be addressed from a single unified perspective. Stochastic control theory offers such a perspective by considering how control strategies should be designed to deal with uncertainty and to benefit from observations [3].

Here we consider an approach to stochastic control that is especially useful for real-time operation of complex spatially distributed systems with uncertain inputs. Such systems are often encountered in geoscience applications, including management of water and petroleum resources, contaminant remediation, and carbon sequestration. Our particular approach to stochastic control uses ensembles of possible inputs

*Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge MA (blin@mit.edu).

†Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge MA (dennism@mit.edu).

and states to represent uncertainty, to determine how to make best use of available measurements, and to select control strategies that are likely to achieve particular objectives.

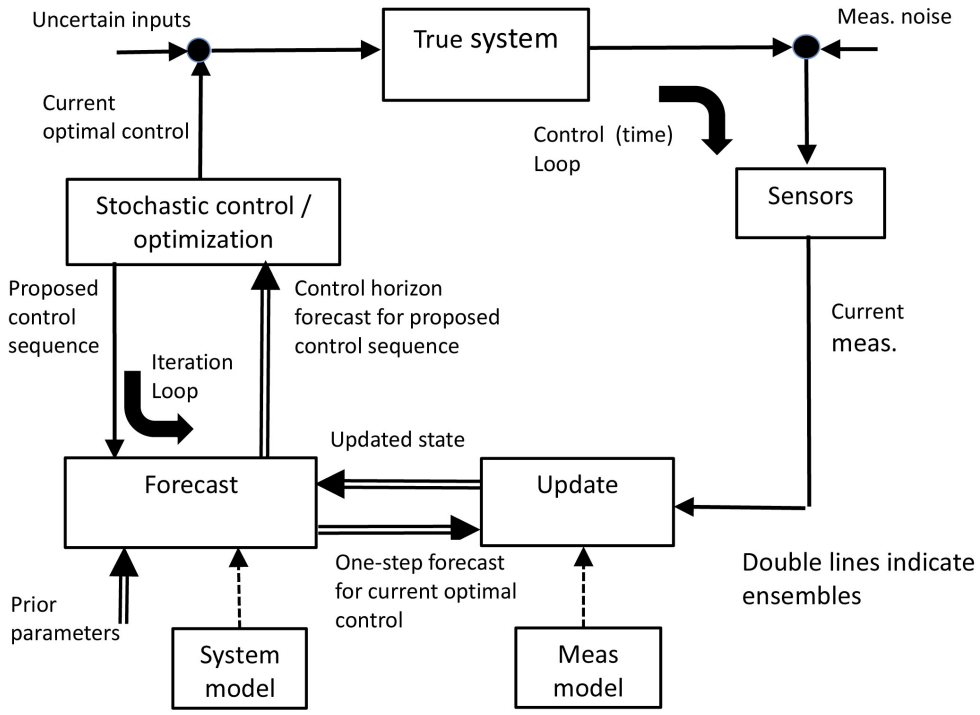

FIG. 1.1. *Basic components of an ensemble feedback control procedure, showing the role of real-time characterization.*

Fig. 1.1 illustrates one possible approach to ensemble real-time control. The true system is subject to uncertain inputs that lead to uncertainties in the system states. We suppose that controlled inputs are applied to this system at discrete times defined over a specified control horizon. Each pass around the control loop in the figure describes operations that take place between two successive control times. We also suppose that noisy measurements of some of the system states are available at the control times. Although these measurements are known up to and including the current time, the states they depend upon are uncertain.

We suppose that the system uncertainties can be described by a prior ensemble of system states composed of many replicates generated from different combinations of possible inputs. In the diagram double lines indicate ensembles while single lines indicate individual variables that may be specified or uncertain, depending on context. The ensemble of possible system states is updated with new measurements whenever they become available. At each control time the current updated ensemble serves as an initial condition for a forecast to the next time and beyond, to the end of the control horizon. Also at each time a single deterministic control is applied to the true system. This control is chosen to minimize a cost-related measure of performance that depends on the entire forecast ensemble so the control explicitly accounts for the fact that the system response is uncertain. A stochastic approach typically leads to

more conservative strategies than would be obtained with a controller that assumes the system response is perfectly known.

The procedure outlined above is a "closed-loop" ensemble controller in the sense that there is feedback from the true system through the measurements and forecast to the controller. This is in contrast to an "open loop" ensemble controller in which there is no feedback either because measurements are not collected or because they have no impact on subsequent controls. Fig. 1.1 is reduced to open loop ensemble control when the measurement update operation is omitted. We can expect closed-loop control to work better than open-loop control when the true system behavior is uncertain since the closed-loop controller can continually adjust as the true system responds, sometimes in unexpected ways [13].

Candidate updating techniques suitable for ensemble real-time control include various forms of the ensemble Kalman filter [7, 8] and the particle filter [9]. These techniques are appropriate for many nonlinear problems and also offer considerable flexibility for describing prior uncertainty. Here we compute the replicates needed to derive real-time controls from an unbiased square root version of the ensemble Kalman filter [11]. There are many methods for deriving the control strategy from the forecast ensemble. We use sequential quadratic programming based on gradients computed with an adjoint approach [12].

As indicated in Fig. 1.1, a system model is used to generate forecast ensembles required by both the updating and optimization steps of the real-time control algorithm. The number of model simulations required is on the order of the number of ensemble replicates multiplied by the number of optimization search iterations, multiplied by the number of control times. Generally speaking, larger problems require more replicates and more search iterations as well as more computation time for each model simulation. The computational effort associated with these simulations can be substantial, to the point where it discourages use of an ensemble approach or at least greatly reduces the size of the ensembles used to guide control decisions. It is difficult to properly describe uncertainty and to derive accurate updates when the ensemble size becomes too small. In this paper we address the computational limitations of ensemble control by using a low-dimensional reduced-order model to make forecasts for both update and optimization operations. The reduced-order model is derived from a higher-dimensional full order model that accurately solves the governing equation used to describe the true system. The model reduction process is designed to preserve the key dynamical properties of the full order model and also to reproduce the probabilistic characteristics of ensembles obtained from full order simulations.

The ensemble control and reduced-order modeling concepts introduced here are tested on a numerical experiment inspired by a practical environmental control problem: cost-effective removal of dissolved contaminants from a groundwater aquifer. The controls of interest in this problem are well pumping rates and the primary sources of uncertainty are poorly known geological properties. The uncertain aquifer properties make the system's response control actions uncertain. Our numerical experiment investigates, for this example, performance differences between open and closed-loop control, the feasibility of using a reduced-order model, and the merits of minimizing mean cost vs. a weighted performance measure that also penalizes cost uncertainty. Our discussion begins with a formulation of the problem based on a full order forecast model and then considers a revised formulation that relies on a reduced-order approximation. We then describe the numerical experiment and present some typical results. We conclude with a consideration of limitations and possible enhancements

of the approach.

2. Problem formulation and methodology.

2.1. Full order formulation. We formulate the real-time control problem by following Fig. 1.1 from the true system through measurement collection and updating, to forecasting and control, back to the true system. The spatially distributed system to be controlled is described by a set of partial differential equations discretized over an appropriate spatial grid. The discretization produces the following set of nonlinear state equations:

$$\frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x}, \alpha, \omega, \mathbf{u}), \quad \mathbf{x}(t_0) = \mathbf{x}_0 \text{ specified}, \tag{2.1}$$

$$\mathbf{G}(\mathbf{x}, \alpha, \omega, \mathbf{u}) = \mathbf{0}. \tag{2.2}$$

Here we have explicitly distinguished two types of spatially discretized equations, both derived from the original partial differential equations. These are a set of ordinary differential equations (2.1) derived from equations with both time and space derivatives and a set of algebraic equations (2.2) derived from equations with only space derivatives. The algebraic functions $\mathbf{F}(\mathbf{x}, \alpha, \omega, \mathbf{u})$ and $\mathbf{G}(\mathbf{x}, \alpha, \omega, \mathbf{u})$ are generated by the discretization procedure and account for the effects of known boundary conditions. In the state equations we distinguish an $N_{\mathbf{x}}$ dimensional vector of dynamic spatially discretized dependent variables (or states) $\mathbf{x}$. It is convenient for control purposes to divide the system inputs into an N_α dimensional vector of time-invariant but spatially variable model parameters α, an $N_\omega$dimensional vector of time-varying uncontrolled inputs ω, and an $N_{\mathbf{u}}$ dimensional vector $\mathbf{u}$ of time-varying control variables that are adjusted to achieve a specified objective. In some cases, it is useful to include the initial condition $\mathbf{x}_0$ in the parameter vector.

In the applications of interest here we suppose that the parameter α is the only uncertain model input. The uncontrolled time-varying input ω is assumed to be perfectly known and the control $\mathbf{u}$ is to be determined. The functions $\mathbf{F}$ and $\mathbf{G}$ used to describe the system are also assumed to be perfectly known. This problem formulation is sometimes called a perfect model scenario. In the real-time control example considered here parameter errors dominate and the perfect model assumption is a reasonable approximation. In other applications it may be necessary to explicitly consider uncertainties in ω. This extension is feasible within the conceptual framework presented here, although it complicates the computations.

It is often convenient to work with an $N_y = N_x + N_\alpha$ dimensional augmented state vector $\mathbf{y}$ that combines dynamic states and time-invariant parameters:

$$\mathbf{y} = \begin{bmatrix} \mathbf{x} \\ \alpha \end{bmatrix}$$

. This vector is described by the following augmented state equations:

$$\frac{d\mathbf{y}}{dt} = \begin{bmatrix} \mathbf{F}(\mathbf{x}, \alpha, \omega, \mathbf{u}) \\ \mathbf{0} \end{bmatrix} = \mathbf{F}_a(\mathbf{y}, \omega, \mathbf{u}), \quad \mathbf{y}(t_0) = \mathbf{y}_0 \text{ specified}, \tag{2.3}$$

$$\mathbf{G}_a(\mathbf{y}, \omega, \mathbf{u}) = \mathbf{G}(\mathbf{x}, \alpha, \omega, \mathbf{u}) = \mathbf{0}. \tag{2.4}$$

Note that (2.3) relies on the assumption that α is time-invariant.

We define the control horizon to be the time interval that extends from the current time $t = k\Delta t$ to a specified end time $T = N_T\Delta t$, where $k = 0, \ldots, N_T - 1$. For simplicity, we suppose that this horizon is divided into $N_T - k$ control intervals of equal duration Δt. We seek a set of controls $\mathbf{u}_{k:N_T-1} = [\mathbf{u}_k, \mathbf{u}_{k+1}, \ldots, \mathbf{u}_{N_T-1}]$ to be applied over the $N_T - k$ remaining control intervals, where each $\mathbf{u}_l$ is an $N_\mathbf{u}$ dimensional vector of scalar controls. We use the colon notation (such as $k : N_T - 1$ in $\mathbf{u}_{k:N_T-1}$) as shorthand for a sequence of variables defined at different times. The control sequence is constrained to be piecewise constant, with each control value held constant within its respective control interval.

We suppose here that the measurements for closed-loop control are obtained at the times $t_{1:N_T-1}$ (i.e. at the beginning of each control period after the first). For simplicity, we also suppose that these measurements are linearly related to the states as follows:

$$\mathbf{z}_k = \mathbf{M}_k\mathbf{y}_k + \mathbf{v}_k, \tag{2.5}$$

where $\mathbf{z}_k$ is an $N_\mathbf{z}$ dimensional measurement vector, $\mathbf{v}_k$ is a zero-mean temporally uncorrelated $N_\mathbf{z}$ dimensional random noise vector with a specified covariance matrix $\mathbf{R}_k$ and $\mathbf{M}_k$ is a specified $N_\mathbf{z}$ by $N_\mathbf{y}$ dimensional measurement matrix.

A stochastic approach to real-time control accounts, at least in an approximate way, for uncertainties about the system to be controlled. In order to do this we adopt a Bayesian approach and suppose that the uncertain augmented state vector is characterized by a randomly generated prior probability distribution or, alternatively, by an ensemble of samples (or replicates) drawn from that distribution. This ensemble is updated whenever measurements become available. As mentioned in the Introduction, our updated ensemble is obtained from an unbiased square root version of the ensemble Kalman filter [11] and our forecast ensemble is obtained from a simulation model based on (2.3) and (2.4) and initialized with the updated ensemble. The net result is an ensemble of conditional forecasts $\mathbf{y}^j_{k:N_T|k}$, indexed by replicate $j = 1, \ldots, N_r$, that describe the system's response over the control horizon $[k\Delta t : N_T\Delta t]$, for a specified control sequence $\mathbf{u}_{k:N_T-1}$. Here the vertical bars indicate conditioning on measurements (e.g. the replicate $\mathbf{y}^j_{k+1|k}$ is a sample from a probability distribution conditioned on measurements collected through $k\Delta t$). The detailed computations required to obtain these forecasts are described below.

In the version of real-time control described here the effectiveness of a proposed control strategy $\mathbf{u}_{k:N_T-1}$ is measured by the cost J_k of operating the system from the current time $k\Delta t$ to the end time $N_T\Delta t$ (i.e. the cost-to-go), for $k = 0, \ldots, N_T - 1$. Costs before the current time have already been incurred and are not considered in the selection of current and future controls. We suppose that the cost-to-go is separable so it can be written as the sum of the costs incurred over the remaining control intervals:

$$J_k(\mathbf{u}_{k:N_T-1}, \mathbf{y}_{k:N_T|k}) = \sum_{l=k}^{N_T-1} \gamma_{l|k}(\mathbf{y}_{l|k}, \mathbf{y}_{l+1|k}, \mathbf{u}_l), \tag{2.6}$$

where $\gamma_{l|k}(\mathbf{y}_{l|k}, \mathbf{y}_{l+1|k}, \mathbf{u}_l)$ is the cost incurred over control interval $[l\Delta t, (l+1)\Delta t]$, conditioned on measurements collected through $k\Delta t$. Note that $\gamma_{l|k}$ is uncertain, for a given control sequence, by virtue of its dependence on the uncertain conditional states $\mathbf{y}_{k:N_T|k}$. We can account for this uncertainty by constructing, for a given control sequence, an ensemble of cost-to-go replicates $J^j_k(\mathbf{u}_{k:N_T-1}, \mathbf{y}^j_{k:N_T|k})$ from an ensemble

$\mathbf{y}^j_{k:N_T|k}$ of conditional states obtained, for example, from an ensemble Kalman filter:

$$J^j_k(\mathbf{u}_{k:N_T-1}, \mathbf{y}^j_{k:N_T|k}) = \sum_{l=k}^{N_T-1} \gamma^j_{l|k}(\mathbf{y}^j_{l|k}, \mathbf{y}^j_{l+1|k}, \mathbf{u}_l). \tag{2.7}$$

These replicates can be used to estimate the probability density of the cost-to-go obtained for a particular control sequence.

Equation (2.7) reveals an assumption implicit in our definition of cost. In this equation the ensemble $\gamma^j_{l|k}$ of costs incurred in control interval $[l\Delta t, (l+1)\Delta t]$ is conditioned only on measurements collected through $k\Delta t$. However, when this cost is actually incurred measurements will have been collected through $l\Delta t$ and the uncertainty will be less than suggested by the $\gamma^j_{l|k}$ ensemble of (2.7). Derivation of a control based on (2.7) essentially ignores the value of information gained from updates with measurements collected in the future. Although these measurements are unknown at time $k\Delta t$ we do know that they will add information, as reflected in the measurement equation of (2.5). We can contrast a stochastic control approach based on (2.7) to stochastic dynamic programming, which accounts for the information obtained from all measurements, past, present, and future, when deriving a real-time decision rule [3]. Although a control algorithm based on (2.7) is suboptimal compared to stochastic dynamic programming, it is much less computationally demanding, especially for large distributed parameter problems.

With the cost-to-go ensemble specified we can consider how to select a single control $\mathbf{u}_k$ to apply to the true system at the current time. Although the current state and the cost-to-go are random the current control must be deterministic. One option for obtaining this control is to minimize the sample mean of the cost-to-go, obtained by averaging over all replicates in the cost-to-go ensemble. A more general approach is to minimize a weighted combination of the sample mean and variance, which can be expressed as:

$$V_k(\mathbf{u}_{k:N_T-1}) = \hat{E}^2\left[J^{1:N_R}_k(\mathbf{u}_{k:N_T-1}, \mathbf{y}^{1:N_R}_{k:N_T|k})\right] + \mu\widehat{\mathrm{Var}}\left[J^{1:N_R}_k(\mathbf{u}_{k:N_T-1}, \mathbf{y}^{1:N_R}_{k:N_T|k})\right], \tag{2.8}$$

where the coefficient μ determines the relative weight given to the squared sample mean $\hat{E}^2$ and the sample variance $\widehat{\mathrm{Var}}$ and the superscript $1:N_R$ emphasizes that the sample moments depend on sums over all N_R state and cost-to-go replicates. As μ increases more emphasis is placed on reducing the spread of the cost-to-go replicates, giving a more robust controller with more predictable results. Robust controllers based on large μ values tend to have larger mean values. In such cases, robustness is achieved at the expense of a higher average cost. Minimization of the deterministic performance function V_k at time $k\Delta t$ yields a set $\mathbf{u}^*_{k:N_T-1|k}$ of optimal controls for the entire control horizon.

We can now combine the minimization of (2.8) at each time with the computation of the required forecasts and updates in a real-time recursion that specifies all required computations. This recursion is conveniently expressed as follows:

At each time $t = k\Delta t$, for $k = 0, \ldots, N_T - 1$, carry out an iterative optimization and forecast procedure and then an update, for each replicate $j = 1, \ldots, N_r$, as follows:

Optimization/Forecast: Perform an iterative search for the optimal control sequence $\mathbf{u}^*_{k:N_T-1|k}$ defined by:

$$\mathbf{u}^*_{k:N_T-1|k} = \mathrm{argmin}\left\{V_k(\mathbf{u}_{k:N_T-1})\right\}, \tag{2.9}$$

subject to the following constraints:

$$\mathbf{u}^L \leq \mathbf{u}_{k:N_T-1} \leq \mathbf{u}^U,$$

where $\mathbf{u}^L$ and $\mathbf{u}^U$ are, respectively, lower and upper bounds on the control variables. The conditional states $\mathbf{y}^j_{k:N_T|k}$ needed to evaluate $J^j_k(\mathbf{u}_{k:N_T-1}, \mathbf{y}^j_{k:N_T|k})$ are related to the proposed control $\mathbf{u}_{k:N_T-1}$ by a set of equality constraints (forecast equations):

$$\frac{d\mathbf{y}^j(t)}{dt} = \mathbf{F}_a(\mathbf{y}^j, \omega, \mathbf{u}), \quad \mathbf{y}^j(k\Delta t) = \mathbf{y}^j_{k|k} \text{ from update at } k\Delta t, \tag{2.10}$$

$$\mathbf{G}_a(\mathbf{y}^j_k, \omega, \mathbf{u}) = \mathbf{0}. \tag{2.11}$$

The one-step ahead forecast $\mathbf{y}^j_{k+1|k}$ obtained from (2.10) is used in the subsequent measurement update at $(k+1)\Delta t$. The optimization problem may be solved with a variety of techniques. If a gradient-based method is used it is convenient to derive the gradients from an adjoint procedure, by using Lagrange multipliers to augment $V_k(\mathbf{u}_{k:N_T-1})$ with (2.10) and (2.11).

Update: After the optimization procedure has converged to an optimal $\mathbf{u}^*_{k|k}$ for the current control interval and the corresponding one-step ahead forecast $\mathbf{y}^j_{k+1|k}$ has been computed, update the ensemble average and then the deviations from the average with new measurements taken at the next time, as follows:

$$\bar{\mathbf{y}}_{k+1|k+1} = \bar{\mathbf{y}}_{k+1|k} + \mathbf{K}_{k+1}[\mathbf{z}_{k+1} - \mathbf{M}_{k+1}\bar{\mathbf{y}}_{k+1|k}], \tag{2.12}$$

$$\mathbf{Y}_{k+1|k+1} = \mathbf{Y}_{k+1|k}\mathbf{T}_{k+1}, \tag{2.13}$$

$$\mathbf{y}^j_{k+1|k+1} = \bar{\mathbf{y}}_{k+1|k+1} + \tilde{\mathbf{y}}^j_{k+1|k+1}. \tag{2.14}$$

Here $\bar{\mathbf{y}}_{k+1|k}$ and $\bar{\mathbf{y}}_{k+1|k+1}$ are the arithmetic averages of the forecast and updated augmented state replicates $\mathbf{y}^j_{k+1|k}$ and $\mathbf{y}^j_{k+1|k+1}$. The $\mathbf{y}^j_{k+1|k}$ are obtained from the previous forecast step. The columns of the $N_\mathbf{y}$ by N_{rep} matrices $\mathbf{Y}_{k+1|k}$ and $\mathbf{Y}_{k+1|k+1}$ are composed of the ensemble perturbations $\tilde{\mathbf{y}}^j_{k+1|k} = \mathbf{y}^j_{k+1|k} - \bar{\mathbf{y}}_{k+1|k}$ and $\tilde{\mathbf{y}}^j_{k+1|k+1} = \mathbf{y}^j_{k+1|k+1} - \bar{\mathbf{y}}_{k+1|k+1}$, for $j = 1, \ldots, N_r$. The $N_\mathbf{y}$ by N_{rep} dimensional matrix $\mathbf{K}_{k+1}$ is the Kalman gain and the N_{rep} by N_{rep} matrix $\mathbf{T}_{k+1}$ is a weighting matrix that transforms forecast ensemble perturbations into updated ensemble perturbations. Both of these matrices are computed from $\mathbf{Y}_{k+1|k}$ and the measurement error covariance matrix $\mathbf{R}_{k+1}$ [11].

The recursion given in (2.9) through (2.14) must be initialized at $t = k = 0$ with a set of augmented state replicates $\mathbf{y}^j_{0|0}$ and with an initial control sequence to start the iterative optimization algorithm. Note that there is no measurement and no update at the initial time (the first measurement and update are at $k = 1, t = \Delta t$). The initial time (or prior) replicates divide into parameters $\alpha^j_{0|0}$ sampled from a specified prior parameter distribution and dependent variables $\mathbf{x}^j_{0|0}$ sampled from a specified prior distribution of initial states. The prior replicates need to be carefully generated so that they are consistent with the spatial structure likely to be encountered in a given

application. For this reason, the replicate generation process should be application-specific rather than generic. An example is illustrated in the subsurface transport example discussed in Section 3.

Although the optimization problem of (2.9) provides a sequence $\mathbf{u}^*_{k:N_T-1|k}$ of optimal control values for the entire control horizon, these values are only conditioned on measurements taken through the current time. Also, only the first value $\mathbf{u}^*_{k|k}$ in the sequence is actually needed at the current time. Consequently, we can discard the future controls $\mathbf{u}^*_{k+1:N_T-1|k}$ derived from (2.9) and apply $\mathbf{u}_k^{\text{opt}} = \mathbf{u}^*_{k|k}$ to the true system at $k\Delta t$. The discarded controls for the later times $(k+1)\Delta t, \ldots, (N_T-1)\Delta t$ will eventually be replaced by new optimization solutions derived from new information obtained at these times. This approach is a version of Model Predictive Control (MPC), which is widely used because it performs well, is easy to implement, and is computationally efficient.

The expressions given in this section define all of the operations in the real-time stochastic control loop of Fig. 1.1. In the next section we consider an approximate but computationally efficient reduced-order formulation of this real-time control procedure.

2.2. Reduced-order formulation. The stochastic real-time control algorithm described above requires multiple evaluations of the system equations given in (2.6) through (2.11): approximately one evaluation for each ensemble replicate for each iteration of the optimization search algorithm at each control/measurement time. In most distributed parameter applications, including the solute transport problem examined in this paper, the most computationally demanding part of these evaluations is the solution of the forward model equations given in (2.1) and (2.2). The high cost of model solutions has the practical effect of limiting the ensemble size that can be used in such applications. This, in turn, limits the effectiveness of the approach by compromising the accuracy of the measurement update operation, which relies on sample moments of the augmented state vector, and by compromising the accuracy of the optimization procedure, which relies on sample moments of the cost-to-go.

There is clearly a tradeoff between the accuracy obtained by using a high-dimensional computationally demanding system model with a small ensemble vs. using a lower-dimensional less computationally demanding model with a large ensemble. In order to investigate this tradeoff we examine real-time control performance with the reduced-order system model described in *Lin and McLaughlin* [unpublished manuscript].

Linear model reduction techniques such as principle orthogonal decomposition (POD) have been widely applied to the deterministic control of spatially distributed systems. Examples include [4] and [17]. Several recent studies have proposed more advanced model reduction methods that may be suitable for nonlinear stochastic optimization applications. These include adaptive procedures that improve the model reduction process by successively updating linear approximations [14] and trust-region methods [2]. Most of these model reduction methods require computationally expensive on-line modifications of the reduced-order model. By contrast, our method is implemented entirely off-line. We derive a set of basis vectors for model reduction from full-order model snapshot simulations performed outside the real-time control loop. Linear expansions in these basis vectors are inserted in a second-order approximation to the full-order model to give a robust nonlinear reduced-order model that captures key dynamic properties of the true system. This approach makes it possible to move expensive computations off-line so the real-time control operation can use ensembles of reasonable size. The basic form of the reduced-order model is summarized

below. Further details are provided in *Lin and McLaughlin* [unpublished manuscript].

Our model reduction approach replaces the uncertain model parameter vector α and the state vector $\mathbf{x}$ appearing in (2.1) and (2.2) with linear combinations of specified basis functions [16], as follows:

$$\alpha \approx \mathbf{\Theta}\hat{\alpha}, \tag{2.15}$$

$$x \approx \Phi\hat{x}, \tag{2.16}$$

where $\mathbf{\Theta}$ is an orthonormal N_α by $N_{\hat{\alpha}} << N_\alpha$ matrix whose columns are a selected set of discrete cosine transform (DCT) basis vectors and $\mathbf{\Phi}$ is an orthonormal $N_{\mathbf{x}}$ by $N_{\hat{\mathbf{x}}} << N_{\mathbf{x}}$ matrix whose columns are principle orthogonal decomposition (POD) basis vectors. These expressions may be combined to give a concise approximation for the augmented state vector:

$$\mathbf{y} \approx \mathbf{\Psi}\hat{\mathbf{y}}, \tag{2.17}$$

where:

$$\hat{\mathbf{y}} = \begin{bmatrix} \hat{\mathbf{x}} \\ \hat{\alpha} \end{bmatrix}, \quad \mathbf{\Psi} = \begin{bmatrix} \mathbf{\Phi} & 0 \\ 0 & \mathbf{\Theta} \end{bmatrix}.$$

The DCT basis vectors are standard functions frequently used for image compression. The POD basis vectors are the leading $N_{\hat{\mathbf{x}}}$ singular vectors of a so-called snapshot matrix. The columns of this matrix are $N_{\mathbf{x}}$ dimensional state vectors and sensitivity derivatives obtained from a set of typical full-order model simulations. These vectors are snapshots of the system state taken at selected times and run with selected parameter values and control sequences. *Lin and McLaughlin* [unpublished manuscript] provide further details on the process of constructing a snapshot matrix and deriving associated POD basis functions.

The accuracy of the reduced-order model depends largely on how well the truncated DCT and POD basis vectors can capture dominant sources of variability in the parameters and states. The truncation process for the pre-specified set of DCT basis vectors is relatively straightforward [10]. However, the choice of the simulated states and sensitivity derivatives included in POD snapshot matrix is less obvious. The reduced-order model will perform well only if the off-line snapshot simulations reflect the full range of parameter and control values encountered during real-time operation.

Our method for achieving good reduced model performance is to generate a large number of snapshot simulations for different parameter and control combinations. It is reasonable to choose the snapshot parameter values to be a subset of the prior parameter ensemble, which describes our prior understanding of parameter variability. The choice of appropriate control sequences is less obvious since we don't have a prior description of possible controls. In the example considered in Section 3 we generate a representative range of snapshot control sequences by optimizing the deterministic cost-to-go J_k^j for the each of the replicates included in the snapshot parameter set. Regardless of the particular approach used, we can expect the actual characterization and optimization procedures to yield conditional parameters and optimal controls that are different from those selected for the snapshot simulations. However, we have found that snapshots generated as described above span the range of actual conditions reasonably well, giving acceptable reduced model performance. In practice,

the performance of the reduced-order model should be tested for each application to make sure that its various approximations and assumptions are justified.

Once appropriate DCT and POD basis vectors are specified (2.15) and (2.16) can be used to replace each element of $\mathbf{y}$ with linear combinations of $\hat{\mathbf{y}}$ elements. The $\hat{\mathbf{y}}$ elements can be derived by inserting the expansions of (2.15) and (2.16) into a second-order approximation of the full-order augmented state equations in (2.3) and (2.4). The resulting expressions can be manipulated to give a set of quadratic reduced-order equations for $\hat{\mathbf{y}}$, expressed in terms of the perturbation $\delta\hat{\mathbf{y}} = \hat{\mathbf{y}} - \hat{\mathbf{y}}_{\text{nom}}$ about a specified nominal $\hat{\mathbf{y}}_{\text{nom}}$ [*Lin and McLaughlin*, unpublished manuscript]:

$$\frac{d\delta\hat{\mathbf{y}}}{dt} \approx \hat{\mathbf{F}}_{\hat{\mathbf{y}}}\delta\hat{\mathbf{y}} + \frac{1}{2}\hat{\mathbf{F}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}(\delta\hat{\mathbf{y}} \otimes \delta\hat{\mathbf{y}}), \quad \delta\hat{\mathbf{y}}(t_0) = 0, \tag{2.18}$$

$$\hat{\mathbf{G}}_{\hat{\mathbf{y}}}\delta\hat{\mathbf{y}} + \frac{1}{2}\hat{\mathbf{G}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}(\delta\hat{\mathbf{y}} \otimes \delta\hat{\mathbf{y}}) \approx \mathbf{0}, \tag{2.19}$$

where:

$$\hat{\mathbf{y}}_{\text{nom}} = \begin{bmatrix} \mathbf{\Phi}^T \mathbf{x}_{\text{nom}} \\ \mathbf{\Theta}^T \alpha_{\text{nom}} \end{bmatrix}.$$

Here α_{nom} is a specified nominal parameter vector and $\mathbf{x}_{\text{nom}}$ is the nominal state obtained by solving (2.1) and (2.2) with $\alpha = \alpha_{\text{nom}}$. These are the state equations for the reduced-order model.

The elements of the $N_{\hat{\mathbf{y}}}$ by $N_{\hat{\mathbf{y}}}$ dimensional coefficient matrices $\hat{\mathbf{F}}_{\hat{\mathbf{y}}}$ and $\hat{\mathbf{G}}_{\hat{\mathbf{y}}}$ depend on the first derivatives of $\mathbf{F}_a$ while the $N_{\hat{\mathbf{y}}}$ by $N_{\hat{\mathbf{y}}}^2$ dimensional matrices $\hat{\mathbf{F}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}$ and $\hat{\mathbf{G}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}$ depend on the second derivatives, all evaluated at the specified nominal value $\hat{\mathbf{y}}_{\text{nom}}$. All of these derivatives depend in turn on the basis vector matrices $\mathbf{\Theta}$ and $\mathbf{\Phi}$. The Kronecker product $\delta\hat{\mathbf{y}} \otimes \delta\hat{\mathbf{y}}$ generates an $N_{\hat{\mathbf{y}}}^2$ dimensional vector of products of augmented state vector elements. Equations (2.18) and (2.19) form a reduced-order model for the $N_{\hat{\mathbf{y}}} << N_{\mathbf{y}}$ unknowns in the vector $\hat{\mathbf{y}}$. In the solute transport example considered later in this paper the elements of $\hat{\mathbf{F}}_{\hat{\mathbf{y}}}$ and $\hat{\mathbf{G}}_{\hat{\mathbf{y}}}$ depend linearly on the elements of the nominal state and the elements of $\hat{\mathbf{F}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}$ and $\hat{\mathbf{G}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}$ are constants. All of these coefficient matrices can be computed off-line, outside the real-time control loop.

We can now formulate a reduced-order approximation of the full-order real-time control algorithm by writing the control equations developed in Section 2.1 in terms of the low-dimensional augmented state vector $\hat{\mathbf{y}}$, as described below.

At each time $t = k\Delta t$, for $k = 0, \ldots, N_T - 1$, carry out a forecast, an optimization and then an update, for each replicate $j = 1, \ldots, N_r$, as follows:

Optimization/Forecast: Perform an iterative search for the optimal control sequence $\mathbf{u}^*_{k:N_T-1|k}$ defined by:

$$\mathbf{u}^*_{k:N_T-1|k} = \text{argmin}\left\{V_k(\mathbf{u}_{k:N_T-1})\right\}, \tag{2.20}$$

subject to the following constraints:

$$\mathbf{u}^L \leq \mathbf{u}_{k:N_T-1} \leq \mathbf{u}^U.$$

The reduced-order conditional states $\hat{\mathbf{y}}^j_{k:N_T|k}$ needed to evaluate the objective function are related to the proposed control $\mathbf{u}_{k:N_T-1}$ by a set of equality constraints (reduced-order forecast equations):

$$\frac{d\delta\hat{\mathbf{y}}^j}{dt} \approx \hat{\mathbf{F}}_{\hat{\mathbf{y}}}\delta\hat{\mathbf{y}}^j + \frac{1}{2}\hat{\mathbf{F}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}(\delta\hat{\mathbf{y}}^j \otimes \delta\hat{\mathbf{y}}^j), \quad \delta\hat{\mathbf{y}}^j(k\Delta t) = \delta\hat{\mathbf{y}}^j_{k|k} \text{ from update}, \tag{2.21}$$

$$\hat{\mathbf{G}}_{\hat{\mathbf{y}}}\delta\hat{\mathbf{y}}^j + \frac{1}{2}\hat{\mathbf{G}}_{\hat{\mathbf{y}}\hat{\mathbf{y}}}(\delta\hat{\mathbf{y}}^j \otimes \delta\hat{\mathbf{y}}^j) \approx \mathbf{0}, \tag{2.22}$$

$$\hat{\mathbf{y}}^j = \delta\hat{\mathbf{y}}^j + \hat{\mathbf{y}}_{\text{nom}}. \tag{2.23}$$

The one-step ahead forecast $\hat{\mathbf{y}}^j_{k+1|k}$ obtained from (2.21) for the optimal control is used in the subsequent measurement update at $(k+1)\Delta t$. The objective in (2.20) is defined in terms of reduced-order variables as follows:

$$V_k(\mathbf{u}_{k:N_T-1}) = \hat{E}^2\left[J_k^{1:N_R}(\mathbf{u}_{k:N_T-1}, \boldsymbol{\Psi}\hat{\mathbf{y}}^{1:N_R}_{k:N_T|k})\right] + \mu\widehat{\text{Var}}\left[J_k^{1:N_R}(\mathbf{u}_{k:N_T-1}, \boldsymbol{\Psi}\hat{\mathbf{y}}^{1:N_R}_{k:N_T|k})\right]. \tag{2.24}$$

Update: After the optimization procedure has converged to an optimal $\mathbf{u}^*_{k|k}$ for the current control interval and the corresponding one-step ahead forecast $\hat{\mathbf{y}}^j_{k+1|k}$ has been computed, update the ensemble average and then deviations from the average, as follows:

$$\bar{\hat{\mathbf{y}}}_{k+1|k+1} = \bar{\hat{\mathbf{y}}}_{k+1|k} + \mathbf{K}_{k+1}[\mathbf{z}_{k+1} - \mathbf{M}_{k+1}\boldsymbol{\Psi}\bar{\hat{\mathbf{y}}}_{k+1|k}], \tag{2.25}$$

$$\hat{\mathbf{Y}}_{k+1|k+1} = \hat{\mathbf{Y}}_{k+1|k}\hat{\mathbf{T}}_{k+1}, \tag{2.26}$$

$$\hat{\mathbf{y}}^j_{k+1|k+1} = \bar{\hat{\mathbf{y}}}_{k+1|k+1} + \bar{\hat{\mathbf{y}}}^j_{k+1|k+1}, \tag{2.27}$$

where all reduced-order variables (with hats) are defined in the same ways as their counterparts in Section 2.1.

This recursion is initialized at $t = k = 0$ with an ensemble of reduced-order state replicates $\hat{\mathbf{y}}^j_{0|0} = \boldsymbol{\Psi}^T\mathbf{y}^j_{0|0}$ and with an initial control sequence to start the iterative optimization algorithm. Also, a model predictive control approach is used so $\mathbf{u}^{\text{opt}}_k = \mathbf{u}^*_{k|k}$ and $\mathbf{u}^*_{k+1:N_T-1|k}$. The reduced-order optimization, forecasting, and update constraint equations are formulated in terms in terms of the low-dimensional vector $\hat{\mathbf{y}}$ rather than the high-dimensional vector $\mathbf{y}$. This provides an important computational advantage that becomes even more significant for as the system dimension increases.

3. Numerical experiment. The purpose of our numerical experiment is to test the performance of a stochastic real-time control algorithm under controlled conditions. We adopt an approach that accounts for uncertainty in system parameters and measurements but assumes that the equations used to describe the system and uncontrolled system inputs are perfectly known (i.e. we adopt a perfect model assumption). The numerical experiment described here documents performance for a virtual true system obtained by using a typical set of spatially variable parameters. The control strategy is based on a series of virtual measurements that are derived by adding random noise to the true states. In additional experiments not presented here we have looked at other true systems having the same basic structure as our example and have obtained similar results.

In the discussion that follows we are particularly interested in 1) the merits of using closed-loop control (with measurement updates) vs. open loop control (without measurement updates), 2) the performance of controllers based on full vs. reduced-order models of the true system, and 3) the tradeoff between mean performance and performance uncertainty. Our investigation relies on a subsurface solute transport (groundwater remediation) problem that illustrates many of the key issues encountered in spatially distributed real-time control applications.

3.1. Experimental setup. Our experiment considers control of subsurface flow and transport in a thin isotropic confined aquifer with a constant layer thickness, contained in the rectangular horizontal domain shown in Fig. 3.1. In the absence of control a solute originating along the bottom boundary travels upward towards the water supply well at P4. Controlled pumping at wells P1, P2, P3 pulls the solute away from P4 and out of the system, incurring both pumping and treatment costs. The best control satisfies water quality requirements at the supply well while minimizing the cost of remediation.

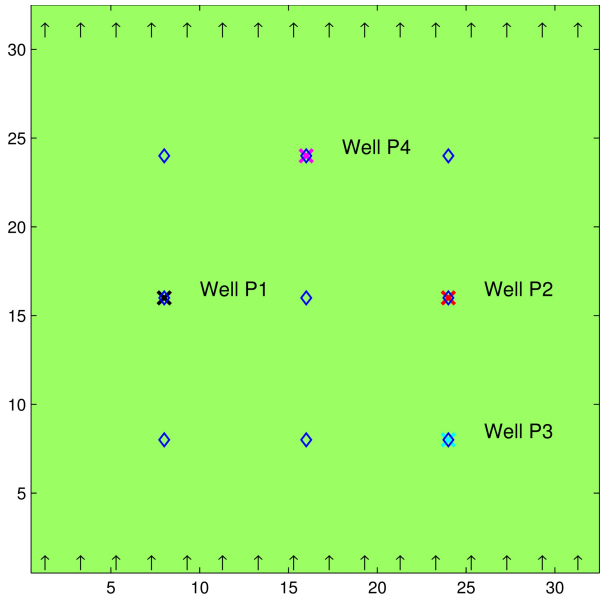

FIG. 3.1. *Numerical experiment simulation domain with* 9 *measurement locations (blue diamonds) and four pumping well locations indicated as* $P1$*,* $P2$*,* $P3$*, and* $P4$*.*

We suppose that measurements of the system response are from fully screened wells and that a two-dimensional vertically-averaged description is appropriate. The corresponding flow and transport model can be written as

$$s\frac{\partial h(r,t)}{\partial t} = -\nabla \cdot q - \sum_{i=1}^{N_w} Q_i \delta(r - r_i), \tag{3.1}$$

$$q(r,t) = -K(r)d\nabla h(r,t), \tag{3.2}$$

$$\theta\frac{\partial c(r,t)}{\partial t} = -\nabla \cdot (qc) + \theta D\nabla^2 c - \sum_{i=1}^{N_w} Q_i c\delta(r - r_i). \tag{3.3}$$

Here $h(r,t)$ is the piezomeric head, $q(r,t)$ is the two-dimensional Darcy velocity, $c(r,t)$ is the solute concentration. s is the specific storage, $Q_i(t)$ is the volumetric flux pumped from the aquifer at well i, $\delta(r - r_i)$ is a two-dimensional spatial Dirac delta function, $K(r)$ is the isotropic hydraulic conductivity, d is the constant aquifer thickness, and θ is the porosity. To simplify the transport equation we assume the

dispersion coefficient D is constant and isotropic in all directions. The head and concentration equations are accompanied by initial and boundary conditions that are discussed in more detail below. Note that the only nonlinearities in these equations are products of the unknown dependent variables h, q_x, q_y, and c, the unknown parameter K, and their respective spatial gradients. That is, these equations are bi-linear.

If we discretize the above model equations over space as described in *Lin and McLaughlin* [unpublished manuscript] the result is:

$$s\frac{\partial \mathbf{h}(t)}{\partial t} - \mathbf{B}\mathbf{q}(t) = \mathbf{R}_1(t), \tag{3.4}$$

$$\mathbf{B}^T\mathbf{h}(t) + \mathbf{A}(\eta)\mathbf{q}(t) = \mathbf{R}_2(t), \tag{3.5}$$

$$\theta\frac{\partial \mathbf{c}(t)}{\partial t} + \mathbf{E}(q)\mathbf{c}(t) + \mathbf{L}\mathbf{c}(t) = \mathbf{R}_3(t), \tag{3.6}$$

where $\mathbf{h}$ is an $N_\mathbf{h}$ dimensional column vector containing the unknown grid cell pressure heads, $\mathbf{q}$ is an $N_\mathbf{q} = N_{qx} + N_{qy}$ dimensional column vector containing the x and y components of the unknown grid cell Darcy velocities, and $\mathbf{c}$ is an $N_\mathbf{c}$ dimensional vector containing unknown grid cell solute concentrations.

The components of the $N_\mathbf{h}$ by $N_\mathbf{q}$ matrix $\mathbf{B}$ are constant while the components of the $N_\mathbf{q}$ by $N_\mathbf{q}$ matrix $\mathbf{A}(\eta)$ depend linearly on the components of the N_α dimensional hydraulic resistivity vector η, which is defined by $\eta_i = 1/K_i$ for each grid cell i. In this application the resistivity is the only parameter included in the generic parameter vector α. Although our model and Kalman filtering algorithms are formulated in terms of resistivity we report results in terms of the more familiar log hydraulic conductivity. The $N_\mathbf{h}$ and $N_\mathbf{q}$ dimensional right-hand-side vectors $\mathbf{R}_1$ and $\mathbf{R}_2$ depend on the flow boundary conditions, including those imposed at the pumping wells. The matrix $\mathbf{E}(\mathbf{q})$ in (3.6) depends nonlinearly on the velocity vector $\mathbf{q}$ when upwinding is used to mitigate numerical dispersion. The matrix $\mathbf{L}$ is constant if the dispersion coefficients are assumed to be constants that do not depend on velocity. If the solute is assumed to exit primarily through the wells and the flow rates at the wells are prescribed the pumping term in (3.6) can be incorporated into the matrix $\mathbf{E}$ and $\mathbf{R}_3 = \mathbf{0}$.

The discretized flow equation system given in (3.4) and (3.5) can be further simplified if we suppose that the pumping rates are constant between control times and that the specific storage is small enough for the flow system to reach an approximate steady state between these times. In this case we only need to solve a steady state version of (3.4) for constant $\mathbf{h}$ and $\mathbf{q}$ over each interval. Then (3.4) through (3.6) can be put in the generic form given in (2.3) and (2.4), with the augmented state vector $\mathbf{y}$ constructed from $\mathbf{h}$, $\mathbf{q}$, $\mathbf{c}$, and η [*Lin and McLaughlin*, unpublished manuscript].

For our numerical example a $320 \times 320 \times 1$ m. computational domain of Fig. 3.1 is discretized into a $32 \times 32 \times 1$ uniform computational grid of $10 \times 10 \times 1$ m. cells. The spatially uniform porosity has a known value of 0.2. The dispersion coefficient is assumed to be an isotropic flow-independent constant over the whole domain. The flow boundary conditions are zero water flux on the east and west sides and constant head boundaries of 10 m and 0 m on the south and north sides. These conditions impose a regional flow from south to north with local fluctuations induced by well pumping and log hydraulic conductivity variations.

The transport boundary conditions are zero solute flux on the east and west sides and zero solute concentration on the south side, except during injection at the

solute source located in the middle of the southern boundary. Solute is injected on this boundary at a constant rate of $5\,\mathrm{m}^3$/day and constant concentration of 50 mg/L for the first 200 days. Outward solute transport is permitted across the northern boundary but the simulation time is limited to keep the solute boundary flux small. The initial concentration is zero everywhere in the domain. These boundary conditions generate a plume of finite extent that travels south to north, gradually moving past the wells but staying largely in the computational domain during the 400-day simulation period. This period is divided into 4 control intervals of 100 days each. In this experiment only wells P1, P2, and P3 are controlled, while well P4 is used for drinking water supply and has a constant pumping rate of $10\,\mathrm{m}^3$/day. Since the dispersion tensor is isotropic and constant and the flow system is at steady-state over each control interval the transport equation is linear, with coefficients that change instantly when the pumping rates change. This simplification facilitates the reduced-order modeling process.

In the full order control algorithm the controlled pumping rates are adjusted to minimize the objective function given in (2.8), which is constructed from the cost-to-go function in (2.7), with the current cost function for replicate j defined as:

$$\gamma_l^j(\mathbf{y}_{l|k}^j, \mathbf{y}_{l+1|k}^j, \mathbf{u}_l) = \sum_{i=1}^{3} u_{l,i}^2 + w \cdot c_{l,4|k}^j, \tag{3.7}$$

where $u_{l,i}$ is the pumping rate at Well $i = 1, 2, 3$ at time $l\Delta t$ and $c_{l,4|k}^j$ is the concentration forecast at time $l\Delta t$, conditioned on measurements taken through $k\Delta t$. This concentration forecast is one of the elements in the full order augmented state vector $y_{l|k}$. The current cost is a weighted combination of the sum of squared pumping rates at the remediation well, which serves as a surrogate for the remediation pumping cost, and the solute concentration at P4. Ideally, both if these quantities should be as small as possible. The lower and upper bounds for the controlled pumping rates (negative values indicate pumping out of system) are:

$$-50\,\mathrm{m}^3/\mathrm{day} \le u_{l,i} \le 0, \quad i = 1, 2, 3.$$

An alternative formulation that gives similar results is to impose solute concentration requirements as inequality constraints and to include only pumping cost in the objective (e.g. by setting $w = 0$). The expressions given here also apply to the reduced-order version of the objective function, with the equation numbers and variables replaced with their reduced-order counterparts wherever applicable.

An ensemble approach to real-time characterization and control relies on prior parameter and initial state ensembles that describe our imperfect knowledge of system properties before measurements are collected. In a subsurface application these prior ensembles should be geologically realistic and appropriate for the site under investigation while also properly conveying, through replicate-to-replicate variability, uncertainty about the true parameters. One flexible option for generating realistic prior replicates is to use multipoint geostatistical techniques based on training images [6]. Training images can be obtained from local experts or constructed from field data. In any case, they are intended to provide useful qualitative information on spatial structures likely to be found in the true system.

The 50 prior resistivity replicates used in our example are generated from a suitable training image with the multi-point algorithm FILTERSIM [18]. We use this

algorithm to define categorical replicates that delineate boundaries between permeable and less permeable geological facies. Resistivity values are assigned to each of the two facies and the resulting binary resistivity images are approximated with a truncated DCT expansion that uses 100 DCT coefficients, about 10% of the number of pixels. The same expansion is used to obtain the basis functions ($\boldsymbol{\Theta}$ columns) for the reduced-order hydraulic resistivity approximation. Since the DCT expansion is approximate the replicates obtained from the retained coefficients are no longer binary but exhibit gradual transitions between higher and lower resistivity areas.

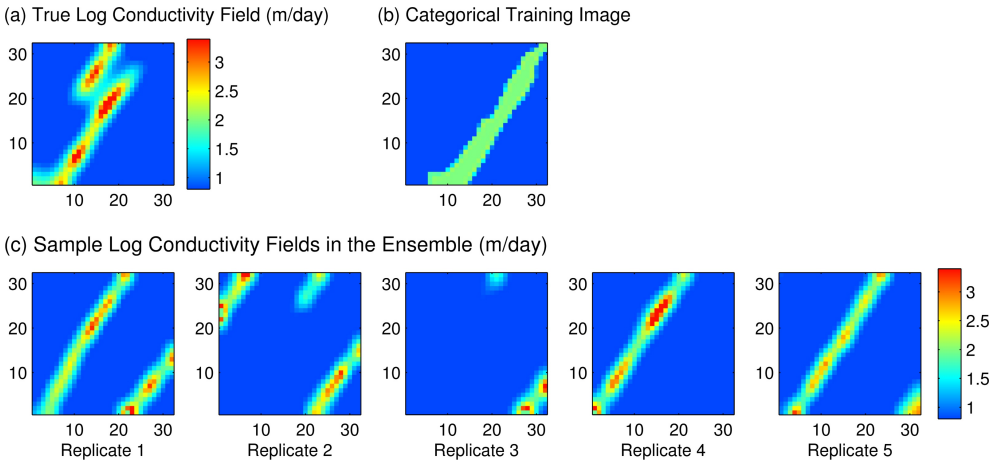

FIG. 3.2. *Some typical log hydraulic conductivity samples. (a) The true log conductivity field; (b) The categorical training image used for generating a prior ensemble in FILTERSIM; (c) Sample prior (unconditional) log conductivity replicates constructed from a truncated DCT approximation to categorical FILTERSIM facies replicates.*

Fig. 3.2(a) shows the single replicate designated as the true field, plotted in log hydraulic conductivity units. Fig. 3.2(b) shows the categorical training image used to generate the true replicate as well as the 50 replicates in the prior ensemble. Some typical members of this ensemble are shown in Fig. 3.2(c). The training image is characterized by a single high-conductivity diagonal channel with a kink in the middle of the domain. This diagonal feature appears, in various degrees, in the true field and in most of the prior replicates. Note that the true replicate is not included in the prior ensemble.

An ensemble Kalman filter is used to update the augmented state replicates with point measurements of head and concentration. These measurements are collected at the nine locations indicated by the blue diamonds in Fig. 3.1. The measurement locations include the four pumping wells. All measurements are taken at the end of the four control steps so there are four updates during the 400-day simulation.

The reduced-order model is constructed from standard DCT basis functions for the model's parameters (resistivities) and customized POD basis functions for the model's dynamic states (heads, velocities, and concentrations). The POD basis functions for each dynamic state (head, two velocity components, and concentration) are obtained from a snapshot matrix composed of full order simulation results for different combinations of possible parameter values and control sequences. Each control sequence minimizes the deterministic costs-to-go for a particular parameter replicate. The head and velocity snapshot outputs include the relevant state as well as deriva-

tives of the state with respect to the 1024 hydraulic resistivities. Each of the resulting head and velocity snapshot matrices has 82320 columns, of which 400 (4 control steps for 100 combinations of 10 replicates and 10 controls) are state vectors and 81920 (4 control steps for 1024 resistivities for 20 combinations of 2 replicates and 10 controls) are derivative vectors. The number of columns in the concentration snapshot matrix is 4000 (40 time steps for 100 combinations of 10 replicates and 10 controls). No derivatives are included in this matrix.

The POD bases are truncated to insure that the retained snapshot matrix singular vectors account for 99%, 93%, 93%, and 99% of the energy in the head, x velocity, y velocity and concentration, respectively. This gives reduced-order dimensions of 56, 103, 97, and 194, for these four variables, each of which originally has a full-order dimension of about 1024. The specific snapshot simulation specifications and POD truncation criteria used here provide a comprehensive picture of full-order model response but could easily be modified, if desired.

3.2. Experimental results. The results for our numerical experiment are organized around several test cases defined as follows:

- Deterministic control: Perfect knowledge of hydraulic conductivity, best possible control performance with no uncertainty,
- Full-order ensemble closed-loop control: Uncertain parameters, measurement updates included,
- Reduced-order ensemble closed-loop control: Uncertain parameters, measurement updates included,
- Full-order ensemble open-loop control: Uncertain parameters, no measurement updates,
- Reduced-order ensemble open-loop control: Uncertain parameters, no measurement updates,
- No control: Reference case with no pumping at control Wells P1, P2, P3. Constant pumping for water supply only at Well P4.

These cases cover a range of control options, varying from a best case with control based on perfect information about the true system to a worst case with no control. The remaining cases represent intermediate options that all account for uncertainty but use different amounts of information and make different approximations. We first examine performance of these cases for $\mu = 0$ in (2.8). This gives a control strategy that minimizes the mean cost-to-go. Then we consider the impact of penalizing the cost-to-go variance by setting $\mu > 0$.

Fig. 3.3 shows the optimal well control sequences and corresponding true cost for each of the controlled cases defined above. Here the true cost is defined as:

$$J_{\text{tot}}(\mathbf{u}^{\text{opt}}_{0:N_T-1}, \mathbf{y}_{0:N_T-1}) = \sum_{l=0}^{N_T-1} \gamma_l(\mathbf{y}_l, \mathbf{y}_{l+1}, \mathbf{u}^{\text{opt}}_l) = \sum_{l=0}^{N_T-1} \left[\sum_{i=1}^{3} \left(u^{\text{opt}}_{l,i}\right)^2 + w \cdot c_{l,4} \right], \tag{3.8}$$

where the state $c_{l,4}$ is the P4 water supply well concentration obtained at time $l\Delta t$ with the true hydraulic resistivity and the actual sequence of optimal controls $\mathbf{u}^{\text{opt}}_{0:N_T-1}$ generated by the real-time algorithm. Fig. 3.4 shows spatial plots of the true solute concentrations obtained by applying the above six control sequences to the true system (simulated with full-order model using the true hydraulic resistivity). Fig. 3.5 shows the true concentration time series at the P4 water supply well for all six sequences. The behavior of the various control options is best understood by examining all three figures.

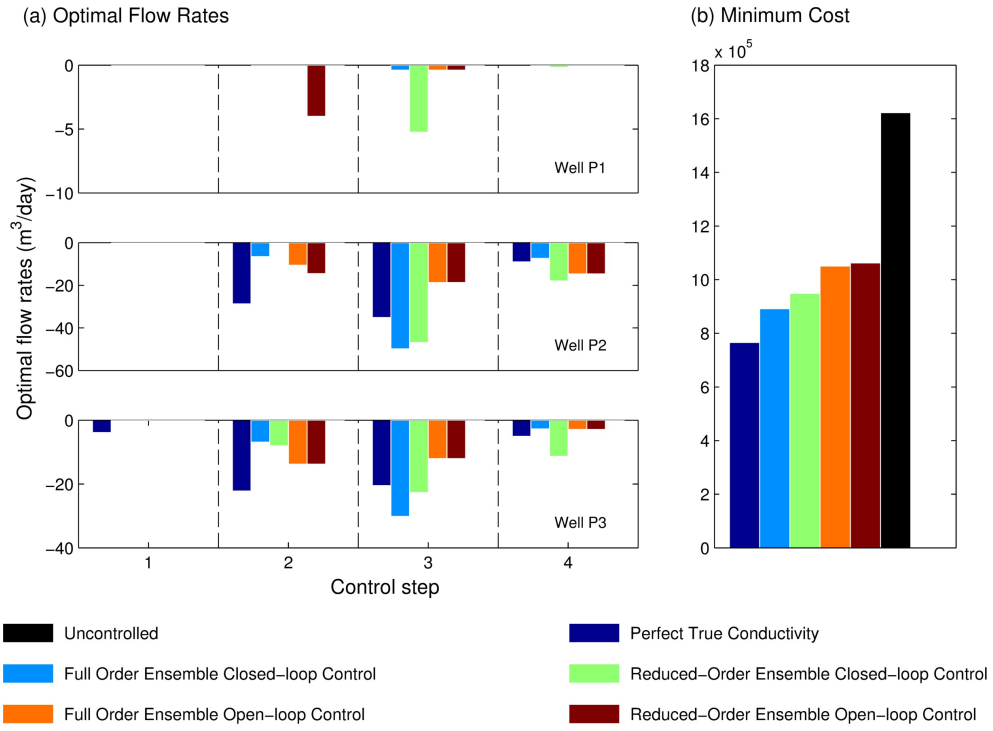

FIG. 3.3. *Comparison of six control strategies: (a) Pumping rates at Wells P1, P2, and P3 in four control intervals; (b) Corresponding true total costs.*

The deterministic control case indicated with dark blue in Fig. 3.3 is characterized by minimal pumping at P1 but significant pumping at P2 and P3 during the second and third control steps. In this case the control strategy allows the natural flow field to move the solute plume to the diagonal permeable channel in the first control interval. Then the solute is extracted through pumping from P2 and P3 during the second and third control intervals. Pumping at these wells is reduced during the fourth control interval since the savings in pumping cost outweighs the penalty incurred from the small concentration at P4.

The full and reduced-order closed loop controllers shown in cyan and green, respectively, pump mostly at P2 and P3 in the third control interval, lagging somewhat behind the deterministic controller. The most distinctive difference between the full and reduced-order strategies is more pumping from P1 in the third control interval for the reduced-order case. This probably occurs because the reduced-order model ensemble locates the permeable channel somewhat closer to P1. The combination of slightly higher concentration at water supply well P4 and more pumping at P1 gives a reduced-order cost that is somewhat higher than the corresponding full-order cost. It is interesting to note that the deterministic control gives lower total cost but higher late-time concentration at P4 than either of the closed loop controls. The lower cost reflects the fact that the peak deterministic pumping rates are lower than those used in the two closed-loop cases.

The full and reduced-order open loop controllers shown in orange and brown, respectively, do not benefit from measurements and distribute pumping more evenly over time. This increases the actual cost since concentrations at water supply well

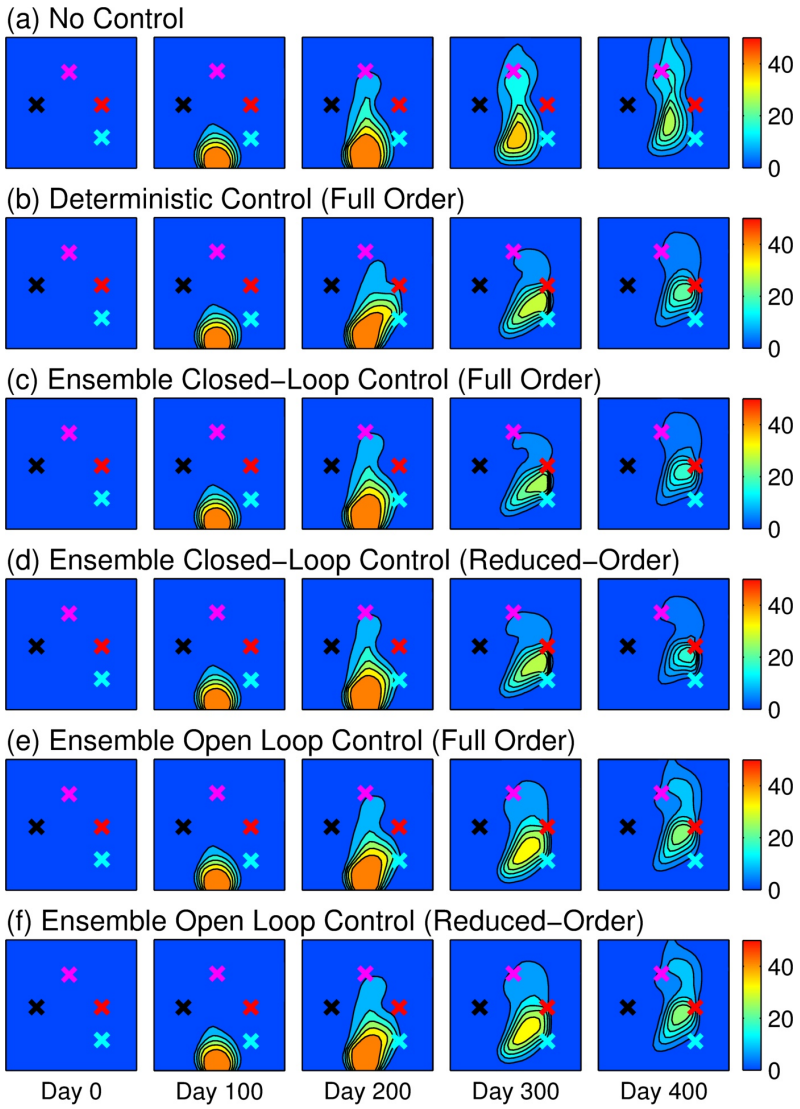

FIG. 3.4. *Spatial distributions of true solute concentration at times* 0, 100, 200, 300, 400 *days for six control strategies with well locations indicated by black (Well P*1*), red (Well P*2*), cyan (Well P*3*), and magenta (Well P*4*) crosses.*

P4 are significantly higher. The open loop controls tend to pump earlier than the corresponding closed-loop controls, suggesting that they predict that the plume will arrive at the permeable channel earlier. The reduced-order control again pumps more at P1 but the overall impact on cost is minor compared to the corresponding full-order case.

Fig. 3.4 reveals that the five controlled solute plumes have similar shapes, However, Fig. 3.5 shows that the closed- and open loop controllers give significantly different concentrations at P4. Overall, the closed-loop controllers work very well, compensating for uncertainty that is not present in the deterministic case. In fact, the performance of the deterministic and closed-loop controllers is quite similar. The cost

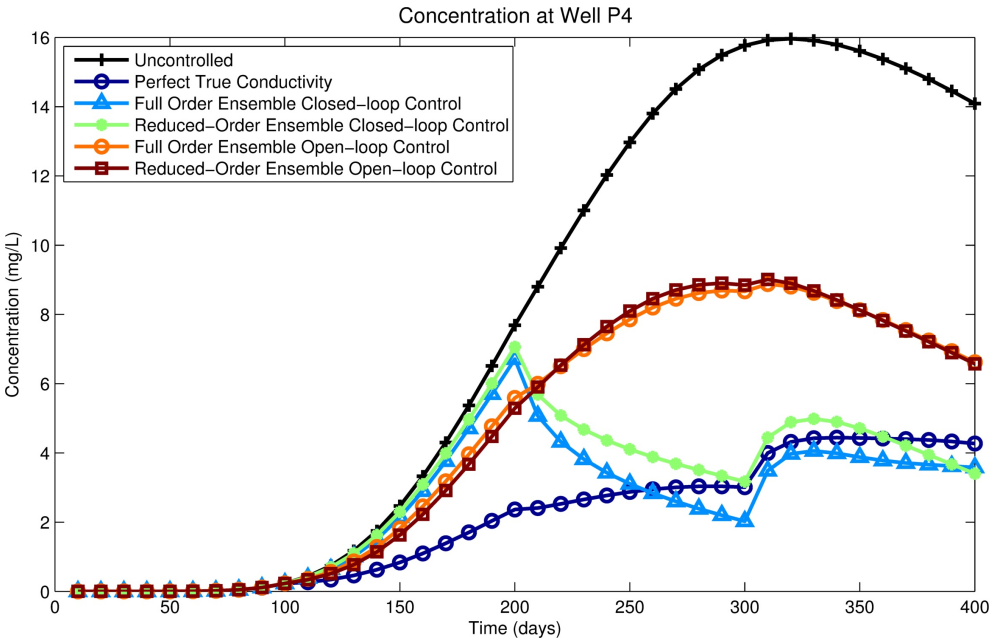

FIG. 3.5. *True concentration time series at Well P4 for six control strategies.*

for the closed-loop cases is about 10% higher than for the deterministic case. This can be viewed as the penalty paid for the moderate level of uncertainty encountered in this particular problem. The discussion below focuses on some of the properties of the closed-loop controllers.

Fig. 3.6 compares the true log hydraulic conductivity to the ensemble means obtained from the full and reduced-order closed loop control algorithms. The mean log conductivity at time 0 is nearly uniform since it averages over many prior replicates with channels placed in different locations (see Fig. 3.2(c)). This scattered representation of geological structure is the one used by the open loop controllers and is the reason these controllers do not perform as well as their closed-loop counterparts, which have access to better geological information as early as the beginning of the second control interval. The well-defined channel apparent in the closed-loop mean log conductivity plots indicates that this channel is in the right location in most of the conditional replicates. Fig. 3.6 also shows that the reduced-order controller gives an ensemble mean log conductivity that is comparable in quality to the full order version.

The control decisions obtained from the closed-loop ensemble algorithms rely on forecasts of solute plume movement. Fig. 3.7 compares the true concentration contours for the two closed-loop controllers to the corresponding ensemble means at various times. Note that the true plumes are different for the full and reduced-order cases because different controls are applied to the true system in these two cases. The figure suggests that the full and reduced-order models give roughly comparable mean characterizations of the corresponding true plumes.

Further insight is provided by Fig. 3.8, which shows time series of the true concentrations and ensembles for each of the four ensemble controllers. The benefits of measurement updating are apparent from a comparison of the open and closed-loop

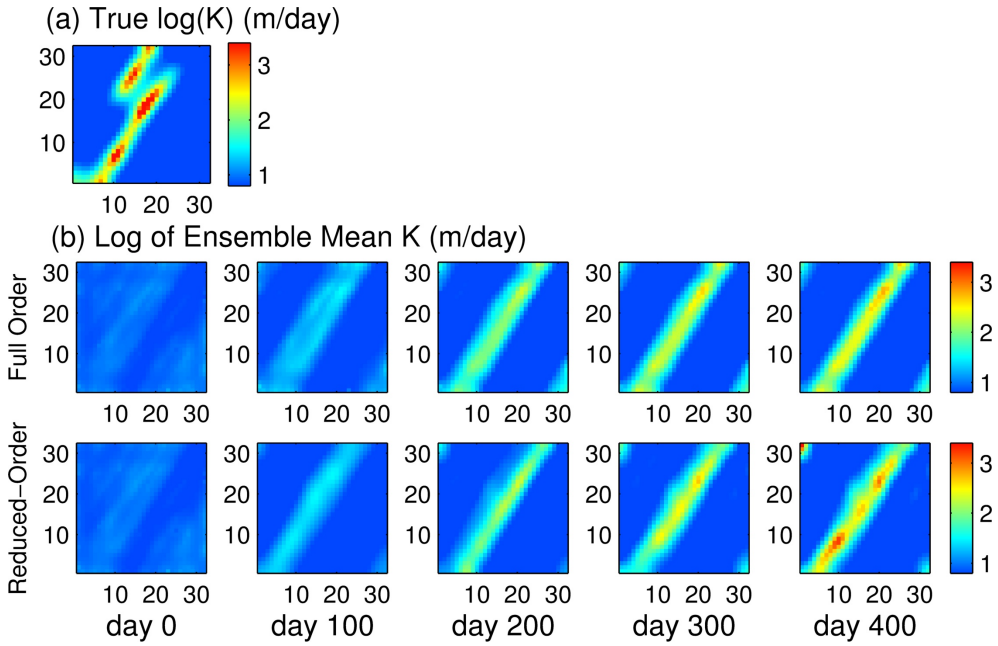

FIG. 3.6. *Spatial distributions of log-hydraulic conductivity: (a) true log-conductivity; (b) log of ensemble mean conductivity for full and reduced-order models.*

concentrations. The closed-loop ensemble means (dashed blue curves) derived from the ensemble (gray curves) closely track the true concentration (red curves) in P2 and P3 and gradually converge to the true in P4. The ensemble spread reflects the EnKF's assessment of uncertainty, which should properly represent the true system uncertainty. When the ensemble collapses the filter ignores new measurement information and its updates can diverge from the true values if the ensemble has not properly captured the true uncertainty. The effects of measurement updating are revealed in Fig. 3.8 by abrupt changes in individual concentration replicates at update times. The measurement updates usually, but not always, bring the replicates closer to the true concentration. The closed-loop full order model gives an ensemble that is somewhat more centered on the true value, with fewer outliers, but the reduced-order model does well overall. The full and reduced-order open-loop results are almost identical, indicating that the reduced-order model is able to capture the system's dynamical behavior and uncertainty over the range of prior resistivity values.

Fig. 3.9 shows the effect of measurement updating on the location and shape of the permeable channel and solute plume. Parts (c) and (d) compare the scattered prior log conductivity replicates to the more clustered conditional replicates obtained after updating. Parts (e) and (f) show a similar comparison for the solute plume. In both cases there is much less variability across the ensemble after updating. The smaller variability that remains reflects the residual uncertainty in log conductivity and concentration.

The results shown in Fig. 3.3 through Fig. 3.9 are all obtained by minimizing the ensemble mean of the cost-to-go. If we also penalize the cost-to-go variance by setting $\mu > 0$ we obtain different control strategies and different system responses for the various ensemble control options. Generally speaking, strategies with $\mu > 0$ are

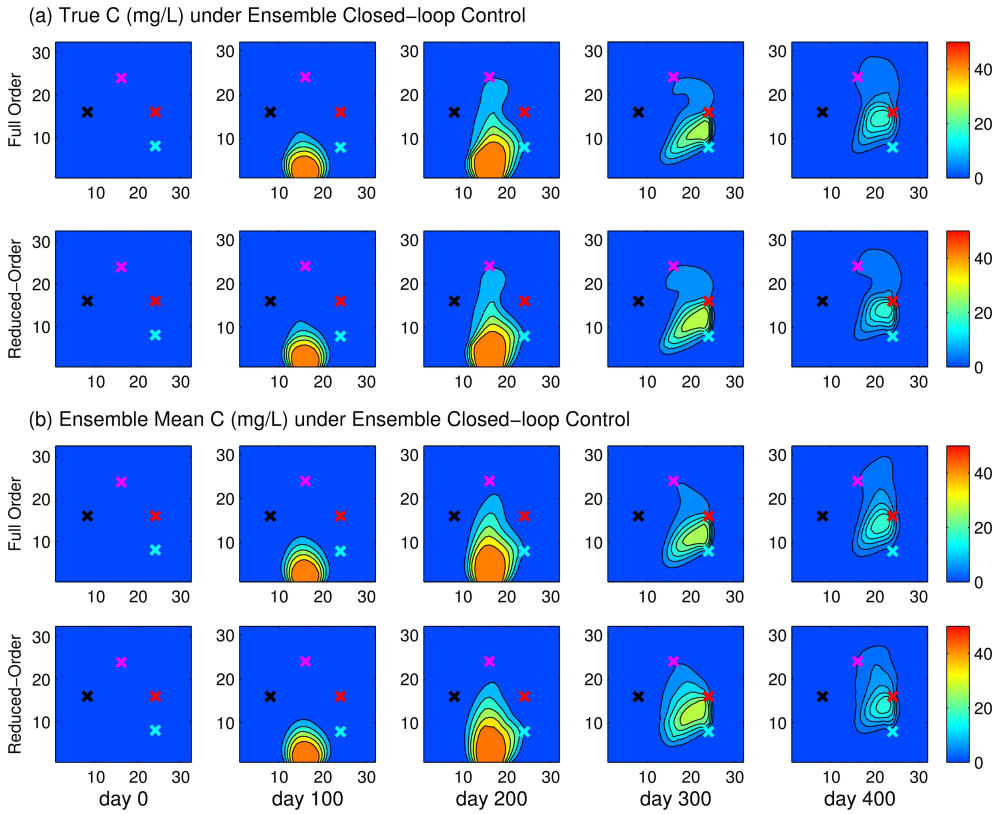

FIG. 3.7. *Spatial distribution of solute concentrations at times* 0, 100, 200, 300, *and* 400 *days with well locations indicated by black (Well P1), red (Well P2), cyan (Well P3), and magenta (Well P4) crosses: (a) true concentration for ensemble closed-loop control with full and reduced-order models with well locations indicated by black (Well P1), red (Well P2), cyan (Well P3), and magenta (Well P4) crosses; (b) ensemble mean concentration for ensemble closed-loop control with full and reduced-order models.*

more conservative, favoring controls that give similar results for different replicates even if the resulting costs are higher than in the $\mu = 0$ case. Fig. 3.10 summarizes the effects of imposing a variance penalty for each of the ensemble control alternatives. The plots in Fig. 3.10(a) and 3.10(c) show probability density functions (PDFs) of the total cost over the entire control horizon, derived from the actual applied open loop control sequence and the prior ensemble of augmented states. This PDF describes our uncertainty about the cost of remediation at the end of the control horizon, after all controls have been derived and applied, with no measurement updating. Each cost replicate used in the PDF computation is defined as:

$$J^j_{\text{open}}(\mathbf{u}^{\text{opt,open}}_{0:N_T-1}, \mathbf{y}^j_{0:N_T|0}) = \sum_{l=0}^{N_T-1} \gamma^j_{l|0}(\mathbf{y}^j_{l|0}, \mathbf{y}^j_{l+1|0}, \mathbf{u}^{\text{opt,open}}_l). \tag{3.9}$$

The plots in Fig. 3.10(b) and 3.10(d) show the PDFs of the total cost over the entire control horizon, derived from the actual applied closed-loop control sequence but using the posterior ensemble of augmented states, conditioned on all measurements collected through the final time. This PDF describes our uncertainty about the cost

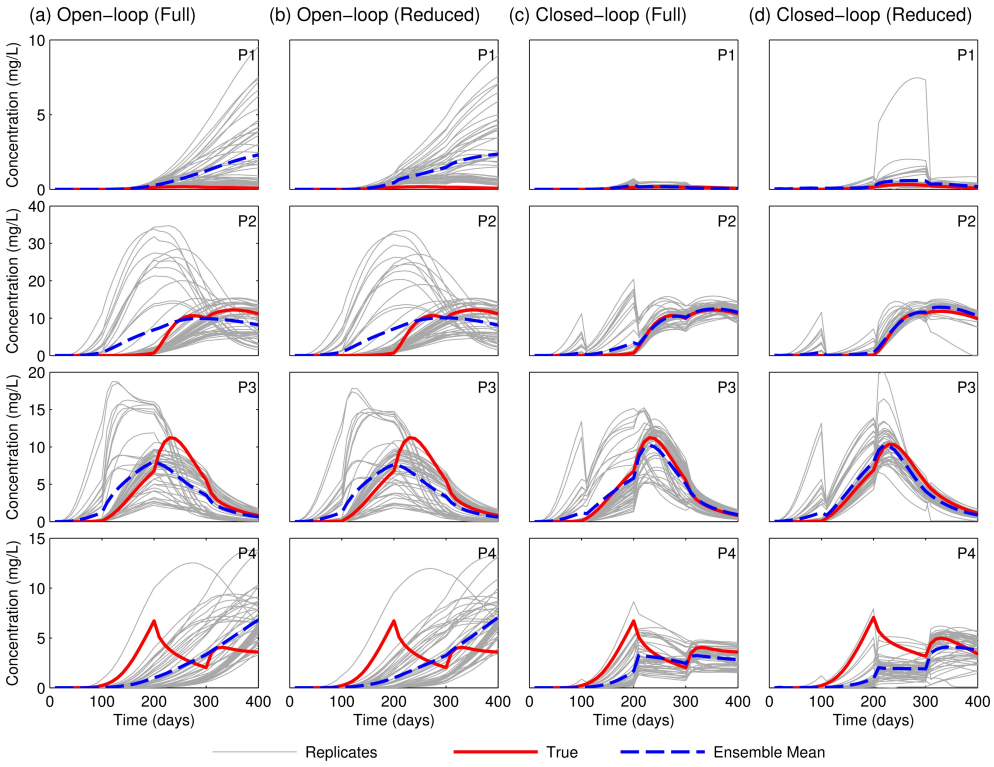

FIG. 3.8. *Concentration time series at the four well locations showing concentration for different log conductivity replicates (gray), concentration obtained with true log conductivity (red), and ensemble mean concentration (blue lines) for: (a) ensemble open loop with full order model; (b) ensemble open loop with reduced-order model; (c) ensemble closed-loop with full order model; (d) ensemble closed-loop with reduced-order model.*

of remediation at the end of the control horizon, after all controls have been derived and applied, with measurement updating. Each cost replicate is defined as:

$$J^j_{\text{closed}}(\mathbf{u}^{\text{opt,closed}}_{0:N_T-1}, \mathbf{y}^j_{0:N_T|N_T}) = \sum_{l=0}^{N_T-1} \gamma^j_{l|N_T}(\mathbf{y}^j_{l|N_T}, \mathbf{y}^j_{l+1|N_T}, \mathbf{u}^{\text{opt,closed}}_l). \tag{3.10}$$

In this case the dynamic states at all times are derived by solving (2.10) and (2.11) forward from the initial time, using the conditional resistivity replicates obtained at the final time. The differences between corresponding open and closed-loop PDFs illustrate the additional information provided by the measurements. Note that Fig. 3.10 shows the distribution of the uncertain cost expended, as derived from known controls and uncertain states. The cost actually billed for the remediation may lie within or outside the range indicated by the PDF since it may be based on other criteria.

The blue curves in Fig. 3.10 show the densities obtained when only the mean is minimized ($\mu = 0$) while the narrower red curves show the densities obtained when the variance is included in the minimization ($\mu = 100$). For reference, the dashed vertical lines show the true cost-to-go for each alternative. It is apparent that the variance penalty significantly reduces the range of costs obtained with different replicates, increasing confidence in the actual cost obtained once the control is applied. The red variance penalty distributions are centered on the corresponding true cost

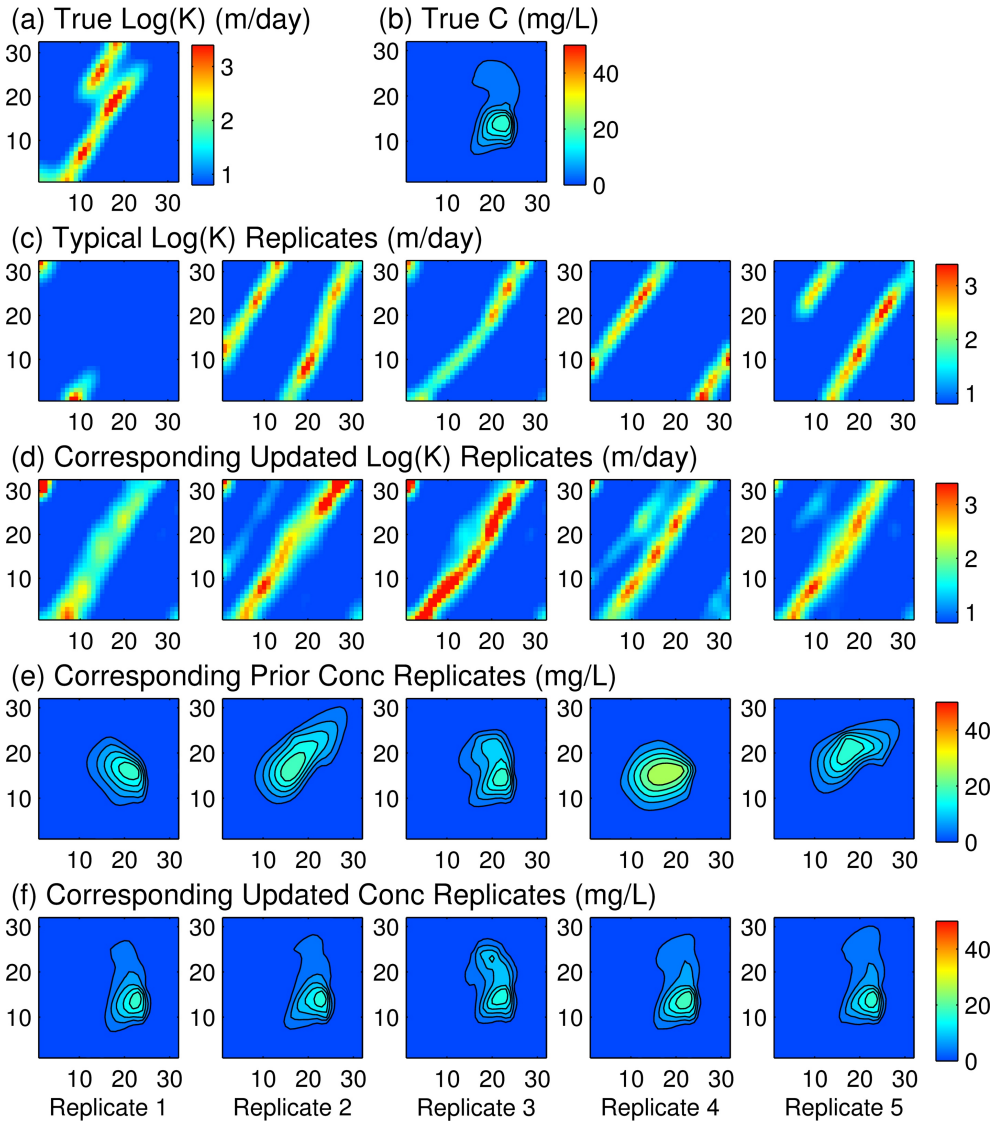

FIG. 3.9. *Spatial distributions for ensemble closed-loop control with reduced-order model: (a) true log-conductivity field; (b) true concentration on day* 400*; (c) five typical prior log-conductivity replicates; (d) five corresponding updated log-conductivity replicates on day* 400*; (e) concentration on day* 400 *for the five prior log-conductivity replicates; (f) corresponding concentration for the five updated log-conductivity replicates.*

values. These true costs are somewhat lower than the true costs obtained without the variance penalty for open loop control but are slightly higher for closed-loop control. The reduced-order model is able to accurately capture the effect of the variance penalty for the open loop case but gives more outliers for the closed-loop case. This probably reflects the fact that the reduced-order POD basis functions are derived from the broader set of prior head and concentration replicates used in open-loop control rather than the narrower set of replicates used in the closed-loop case. Nevertheless, the reduced-order model captures the overall effect of the variance

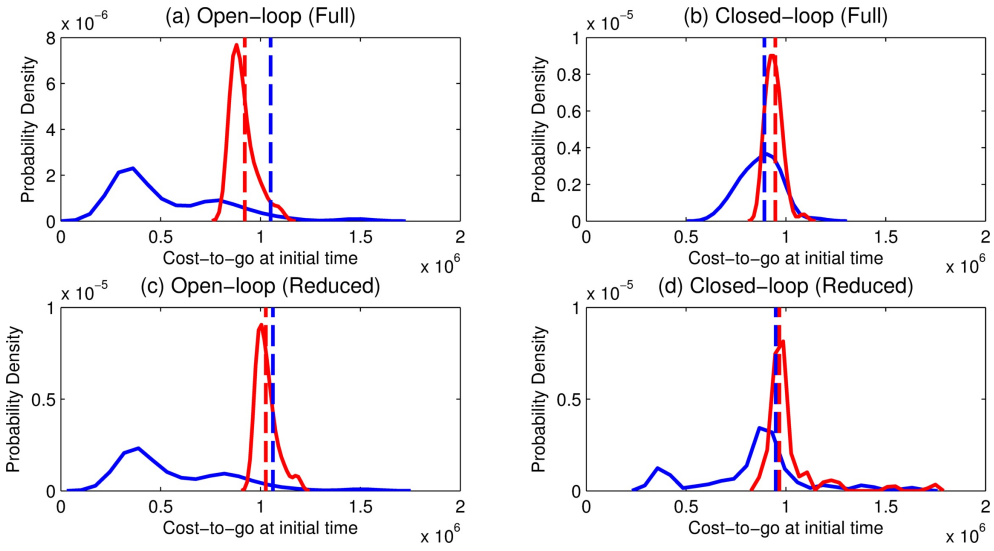

FIG. 3.10. *Probability densities of the total cost for: (a) full order open loop control; (b) full order closed-loop control; (c) open loop reduced-order control; (d) closed-loop reduced-order control. Blue lines indicate* $\mu = 0$ *(without variance penalty), while red lines are for* $\mu = 100$ *(with variance penalty). Vertical dashed lines indicate the cost obtained with the true resistivity.*

penalty on the cost-to-go probability densities for both open and closed loop cases. These results suggest that the variance penalty provides significantly more confidence in control performance for a minimal increase in actual cost, at least for the example considered here.

Overall, the results of our numerical experiment clearly demonstrate the benefits of closed-loop control, which significantly improves performance by correcting for the effects of parameter uncertainty. The results also show that the reduced-order model proposed here is an acceptable substitute for a much more expensive full order model. In our example the reduced-order model achieves control performance comparable to the full order case while requiring only 20% of the computation time. The computational benefit of a reduced-order model will probably be much greater for larger problems. We have restricted our attention here to a full order problem with about 4000 unknowns because of the substantial computational requirements needed to obtain full order results for a performance comparison. The primary cost of reduced-order control is the off-line derivation of the reduced-order coefficient matrices. On-line (real-time) computational costs are a small fraction of the off-line costs. This implies that reduced-order ensemble control is particularly attractive when the time available for on-line computation is limited and it is feasible to move substantial amounts of computation off-line.

4. Discussion and conclusions. The numerical experiment described in this paper demonstrates the feasibility of ensemble closed-loop control with a reduced-order model for a particular solute transport problem. The general concepts presented in Section 2 and applied to the example in Section 3 are broadly applicable. In particular, it is possible to use in many different contexts i) an ensemble of possible states to describe uncertainty, ii) a Kalman filter to update the ensemble with measurements, iii) a reduced-order model to make ensemble forecasts, and iv) a stochastic

model predictive control procedure. The key question is how well these concepts will work in any particular application. Ultimately, this question needs to be answered with problem-specific experiments. However, we have observed in our solute transport example some outcomes that have broad significance.

The basic advantages of closed-loop deterministic control are well known. The advantages of stochastic control are less extensively documented but it seems apparent from our experiments that a stochastic approach provides considerable flexibility to chose controls that account for uncertainty in a systematic way. The tests summarized in Fig. 3.10 indicate that closed-loop controls can be adjusted to provide more predictable results (less uncertainty in cost) with little or no sacrifice in average performance. This general result probably applies in many fields, although the specific nature of the tradeoff between average performance and performance uncertainty is difficult to determine without application-specific tests.

It is more difficult to assess the general feasibility and benefit of using a reduced-order model for ensemble real-time control. The importance of computational issues is related to the control and system response time scales. In general, it is desirable to change controls over time scales comparable to the time it takes for the system to respond to typical changes in uncertain inputs. If the system states change significantly over a control interval it is more difficult to monitor and control performance. On the other hand, if measurements and control decisions are made frequently there is less time to carry out the extensive computations required by an ensemble real-time control algorithm. So the need for reduced-order modeling depends on the system response and control time scales as well as the computational time required to run a full order model. In problems involving relatively slowly responding geological systems, such as the one considered here, the time required to carry out ensemble control computations with a full order model may be significantly less than the time available to make a decision. On the other hand, in problems involving faster-responding atmospheric systems a reduced-order model may be essential. The type of model reduction used here transfers computational effort from on-line to off-line operations. When such a transfer is needed to make the ensemble approach feasible reduced-order modeling should be seriously considered. Otherwise, it may not be necessary.

The decision to use reduced-order modeling also depends on the feasibility of adequately reproducing the dynamic behavior of the true system with a simplified approximation. The reduced-order model used in our example performs well for a particular true system with a dominant high permeability channel but cannot be expected to perform as well over a wider range of conditions unless it is re-derived from a new set of snapshot matrices. The success of this process depends on the conditions selected for the snapshot simulations. If these conditions cover too narrow a range of possibilities the reduced-order model will perform well only in certain limited situations. If the range of snapshot simulations is too wide the reduced-order model may not work particularly well for any condition, even one falling within this range.

The approach to reduced-order modeling described here is only one of many possibilities. We selected it largely because it is well suited for our example problem. For this example the quadratic expansion used in the derivation of Section 2.2 is exact and the resulting reduced model does not depend on the specified nominal state. In other applications it may be preferable to use alternative model reduction techniques such as the Discrete Empirical Interpolation Method (DEIM) or methods that require on-line adjustments to the reduced-order approximation [1, 5, 15]. Although there is

still much room for experimentation, our research suggests that the basic concept of combining ensemble measurement updating, forecasting, and optimization with reduced-order modeling has much merit and can lead to effective control strategies that are able to deal with uncertainty.

Acknowledgments. We thank Ruben Juanes, Youssef M. Marzouk, and Karen Willcox for their comments. We gratefully acknowledge financial support for this work, provided by Shell and Eni S.p.A..

REFERENCES

[1] M. Bergmann, L. Cordier, and J.-P. Brancher, *Optimal rotary control of the cylinder wake using proper orthogonal decomposition reduced-order model*, Physics of Fluids, 17 (2005), pp. 097–101.
[2] M. Bergmann, L. Cordier, and J. P. Braucher, *Drag minimization of the cylinder wake by trust-region proper orthogonal decomposition*, Active Flow Control, 95 (2007), pp. 309–324.
[3] D. P. Bertsekas, *Dynamic programming and stochastic control*, Mathematics in science and engineering, Academic Press, New York, 1976.
[4] G. A. Birnovskii, *On optimal-control of multiphase porous flow in an oil bed*, Ussr Computational Mathematics and Mathematical Physics, 28 (1988), pp. 156–163.
[5] S. Chaturantabut and D. C. Sorensen, *Nonlinear model reduction via discrete empirical interpolation*, SIAM Journal on Scientific Computing, 32 (2010), pp. 2737–2764.
[6] C. V. Deutsch, *Geostatistical reservoir modeling*, Oxford University Press, Oxford; New York, 2002.
[7] G. Evensen, *Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte-Carlo methods to forecast error statistics*, Journal of Geophysical Research-Oceans, 99 (1994), pp. 10143–10162.
[8] ———, *Sampling strategies and square root analysis schemes for the EnKF*, Ocean Dynamics, 54 (2004), pp. 539–560.
[9] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, *Novel approach to nonlinear/non-Gaussian Bayesian state estimation*, IEEE Proceedings F on Radar and Signal Processing, 140 (1993), pp. 107–113.
[10] B. Jafarpour, *Oil reservoir characterization using ensemble data assimilation*, PhD, Massachusetts Institute of Technology, 2008.
[11] D. M. Livings, S. L. Dance, and N. K. Nichols, *Unbiased ensemble square root filters*, Physica D-Nonlinear Phenomena, 237 (2008), pp. 1021–1028.
[12] J. Nocedal and S. J. Wright, *Numerical optimization*, Springer, New York, 2006.
[13] K. J. Åström and R. M. Murray, *Feedback systems: an introduction for scientists and engineers*, Princeton University Press, Princeton, 2008.
[14] S. S. Ravindran, *Reduced-order adaptive controllers for fluid flows using POD*, Journal of Scientific Computing, 15 (2000), pp. 457–478.
[15] M. J. Rewienski, *A trajectory piecewise-linear approach to model order reduction of nonlinear dynamical systems*, PhD, Massachusetts Institute of Technology, 2003.
[16] L. Sirovich, *Turbulence and the dynamics of coherent structures. I: Coherent structures. II: Symmetries and transformations. III: Dynamics and scaling*, Quarterly of Applied Mathematics, 45 (1987), pp. 561–590.
[17] G. A. Virnovsky, *Water flooding strategy design using optimal control theory*, in the 6th European IOR-Symposium, 1987.
[18] T. Zhang, *Filter-based training pattern classification for spatial pattern simulation*, PhD, Standford University, 2006.